

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Scott C. Johnson et al.

Filed: Concurrently Herewith

For: NETWORK CONNECTED COMPUTING SYSTEM INCLUDING NETWORK SWITCH

Serial No.: Unknown

Prior Group Art Unit: 2152

Examiner: Unknown

Atty. Dkt: SURG:157

NUMBER: EL703719648US

I hereby certify that this paper or fee is being deposited with the United States Postal Service "EXPRESS MAIL POST OFFICE TO ADDRESSEE" service, postage prepaid, under 37 CFR 1.10 on the date indicated below and is address to: Assistant Commissioner of Patents, Washington, D.C. 20231.

Mulri J. Murara
Signature

5-31-01
Date

Assistant Commissioner For Patents
Washington, D.C. 20231

Dear Sir:

PRELIMINARY AMENDMENT TO REDUCE FILING FEES

Please amend the application as follows.

In the specification:

The rewritten clean versions of all the specification changes are provided below.
Attached at the end of this paper is an Appendix providing an indication of the changes relative to the prior version of the specification, as now required by Rule 121.

On page 1, replace the paragraph from lines 1-4 with:

-- Patent Application for
NETWORK CONNECTED COMPUTING SYSTEM INCLUDING NETWORK SWITCH

Inventors: Scott C. Johnson, Brian W. Bailey, Rodney S. Canion, Thomas E. Garvens,
Gregory J. Jackson and Roger K. Richter--

On page 1, please replace the paragraph from lines 6-22 with:

This application is a continuation of Application Serial No. 09/797,413, which was filed March 1, 2001 and is entitled "NETWORK CONNECTED COMPUTING SYSTEM", which in turn claims priority from Provisional Application Serial No. 60/187,211, which was filed March 3, 2000 and is entitled "SYSTEM AND APPARATUS FOR INCREASING FILE SERVER BANDWIDTH," from Provisional Application Serial No. 60/246,343, which was filed November 7, 2000 and is entitled "NETWORK CONTENT DELIVERY SYSTEM WITH PEER TO PEER PROCESSING COMPONENTS," from Provisional Application Serial No. 60/246,335, which was filed November 7, 2000 and is entitled "NETWORK SECURITY ACCELERATOR," from Provisional Application Serial No. 60/246,443, which was filed November 7, 2000 and is entitled "METHODS AND SYSTEMS FOR THE ORDER SERIALIZATION OF INFORMATION IN A NETWORK PROCESSING ENVIRONMENT," from Provisional Application Serial No. 60/246,373, which was filed November 7, 2000 and is entitled "INTERPROCESS COMMUNICATIONS WITHIN A NETWORK NODE USING SWITCH FABRIC," from Provisional Application Serial No. 60/246,444, which was filed November 7, 2000 and is entitled "NETWORK TRANSPORT ACCELERATOR," and from Provisional Application Serial No. 60/246,372, which was filed November 7, 2000 and is entitled "SINGLE CHASSIS NETWORK ENDPOINT SYSTEM WITH NETWORK PROCESSOR FOR LOAD BALANCING," the disclosures of each being incorporated herein by reference.

Please replace the paragraph on page 16 from lines 10 -31 and on page 17 from lines 1-3 with the following:

09/797,412 - 09/797,411

--The processing cores of a network processor are typically accompanied by special purpose cores that perform specific tasks, such as fabric interfacing, table lookup, queue management, and buffer management. Network processors typically have their memory management optimized for data movement, and have multiple I/O and memory buses. The programming capability of network processors permit them to be programmed for a variety of tasks, such as load balancing, network protocol processing, network security policies, and QoS/CoS support. These tasks can be tasks that would otherwise be performed by another processor. For example, TCP/IP processing may be performed by a network processor at the front end of an endpoint system. Another type of processing that could be offloaded is execution of network security policies or protocols. A network processor could also be used for load balancing. Network processors used in this manner can be referred to as "network accelerators" because their front end "look ahead" processing can vastly increase network response speeds. Network processors perform look ahead processing by operating at the front end of the network endpoint to process network packets in order to reduce the workload placed upon the remaining endpoint resources. Various uses of network accelerators are described in the following concurrently filed U.S. patent applications: no. 09/797,412, entitled "Network Transport Accelerator," by Bailey et. al; no. 09/797,507, entitled "Single Chassis Network Endpoint System With Network Processor For Load Balancing," by Richter et. al; and no. 09/797,411, entitled "Network Security Accelerator," by Canon et. al; the disclosures of which are all incorporated herein by reference. When utilizing network processors in an endpoint environment it may be advantageous to utilize techniques for order serialization of information, such as for example, as disclosed in concurrently filed U.S. patent application no. 09/797,197, entitled "Methods and Systems For The Order Serialization Of Information In A Network Processing Environment," by Richter et. al, the disclosure of which is incorporated herein by reference.--

Please replace the paragraph on page 23 from lines 12 - 25 with the following:

--Storage management engine 1040 may employ any suitable method for caching data, including simple computational caching algorithms such as random removal (RR), first-in first-out (FIFO), predictive read-ahead, over buffering, etc. algorithms. Other suitable caching

algorithms include those that consider one or more factors in the manipulation of content stored within the cache memory, or which employ multi-level ordering, key based ordering or function based calculation for replacement. In one embodiment, storage management engine may implement a layered multiple LRU (LMLRU) algorithm that uses an integrated block/buffer management structure including at least two layers of a configurable number of multiple LRU queues and a two-dimensional positioning algorithm for data blocks in the memory to reflect the relative priorities of a data block in the memory in terms of both recency and frequency. Such a caching algorithm is described in further detail in concurrently filed U.S. patent application no. 09/797,198, entitled "Systems and Methods for Management of Memory" by Qiu et. al, the disclosure of which is incorporated herein by reference.--

Please replace the paragraph on page 23 from lines 27 - 31 and on page 24 from lines 1 - 5 with the following:

--For increasing delivery efficiency of continuous content, such as streaming multimedia content, storage management engine 1040 may employ caching algorithms that consider the dynamic characteristics of continuous content. Suitable examples include, but are not limited to, interval caching algorithms. In one embodiment, improved caching performance of continuous content may be achieved using an LMLRU caching algorithm that weighs ongoing viewer cache value versus the dynamic time-size cost of maintaining particular content in cache memory. Such a caching algorithm is described in further detail in concurrently filed U.S. patent application no. 09/797,201, entitled "Systems and Methods for Management of Memory in Information Delivery Environments" by Qiu et. al, the disclosure of which is incorporated herein by reference.--

Please replace the paragraph on page 46 from lines 10 - 16 with the following:

--Information gathered by monitoring agents 245 may be employed for a wide variety of purposes including for billing of individual content suppliers and/or users for pro-rata use of one or more resources, resource use analysis and optimization, resource health alarms, *etc.* In addition, monitoring agents may be employed to enable the deterministic delivery of content by

system 200 as described in concurrently filed, co-pending United States patent application number 09/797,200, entitled "Systems and Methods for the Deterministic Management of Information," which is incorporated herein by reference.--

Please replace the paragraph on page 28 from lines 16 - 26 with the following:

--In yet other embodiments, the processing modules may be specialized to specific applications, for example, for processing and delivering HTTP content, processing and delivering RTSP content, or other applications. For example, in such an embodiment an application processing module 1070a and storage processing module 1040a may be specially programmed for processing a first type of request received from a network. In the same system, application processing module 1070b and storage processing module 1040b may be specially programmed to handle a second type of request different from the first type. Routing of requests to the appropriate respective application and/or storage modules may be accomplished using a distributive interconnect and may be controlled by transport and/or interface processing modules as requests are received and processed by these modules using policies set by the system management engine.--

Please replace the paragraph on page 34 from lines 2 - 13 with the following:

--In yet another embodiment, at least two network interface modules 1030a and 1030b may be provided, as illustrated in FIG. 1A. In this embodiment, a first network interface engine 1030a may receive incoming data from a network and pass the data directly to the second network interface engine 1030b for transport back out to the same or different network. For example, in the remote or live broadcast application described above, first network interface engine 1030a may receive content, and second network interface engine 1030b provide the content to the network 1020 to fulfill requests from one or more clients for this content. Peer-to-peer level communication between the two network interface engines allows first network interface engine 1030a to send the content directly to second network interface engine 1030b via distributive interconnect 1080. If necessary, the content may also be routed through transport

processing engine 1050, or through transport processing engine 1050 and application processing engine 1070, in a manner described above.--

Please replace the paragraph on page 34 from lines 23 - 31 with the following:

--The content delivery system 1010 described above is configured in a peer-to-peer manner that allows the various engines and modules to communicate with each other directly as peers through the distributed interconnect. This is contrasted with a traditional server architecture in which there is a main CPU. Furthermore unlike the arbitrated bus of traditional servers, the distributed interconnect 1080 provides a switching means which is not arbitrated and allows multiple simultaneous communications between the various peers. The data and communication flow may by-pass unnecessary peers such as the return of data from the storage management processing engine 1040 directly to the network interface processing engine 1030 as described with reference to FIG. 1B.--

Please replace the paragraph on page 42 from lines 23 - 31 with the following:

--It will be understood with benefit of this disclosure that the hardware embodiment and multiple engine configurations thereof illustrated in FIGS. 1C-1F are exemplary only, and that other hardware embodiments and engine configurations thereof are also possible. It will further be understood that in addition to changing the assignments of individual processing modules to particular processing engines, distributive interconnect 1080 enables the various processing flow paths between individual modules employed in a particular engine configuration in a manner as described in relation to FIG. 1B. Thus, for any given hardware embodiment and processing engine configuration, a number of different processing flow paths may be employed so as to optimize system performance to suit the needs of particular system applications.--

Please replace the paragraph on page 44 from lines 15 - 28 with the following:

.

--FIG. 2 illustrates a single networking subsystem processor module 205 that is provided to perform the combined functions of network interface processing engine 1030 and transport processing engine 1050 of FIG. 1A. Communication and content delivery between network 260 and networking subsystem processor module 205 are made through network connection 270. For certain applications, the functions of network interface processing engine 1030 and transport processing engine 1050 of FIG. 1A may be so combined into a single module 205 of FIG. 2 in order to reduce the level of communication and data traffic handled by communications path 230 and data movement path 235 (or single switch fabric), without adversely impacting the resources of application processing engine or subsystem module. If such a modification were made to the system of FIG. 1A, content requests may be passed directly from the combined interface/transport engine to network application processing engine 1070 via distributive interconnect 1080. Thus, as previously described the functions of two or more separate content delivery system engines may be combined as desired (*e.g.*, in a single module or in multiple modules of a single processing blade), for example, to achieve advantages in efficiency or cost.--

In the claims:

Please cancel claims 1-132 and add the following new claims:

- 133. A method of configuring a network switch, comprising:
- providing a plurality of separate processor engines, the processor engines being assigned separate tasks in an asymmetrical multi-processor configuration;
 - providing an interface connection to at least one of the processor engines to couple the network switch system to a network; and
 - generating an accelerated data flow through the network switch.
134. The method of claim 133, wherein the separate processor engines communicate as peers in a peer to peer environment.

135. The method of claim 134, wherein the separate processors communicate across a distributed interconnect.
136. The method of claim 135, wherein the distributed interconnect is a switch fabric.
137. The method of claim 135, wherein the processor engine coupling the network switch to a network comprises a network processor.
138. The method of claim 137, further comprising performing look ahead processing within the network processor to off-load processing tasks from the other processor engines.
139. The method of claim 138, wherein the network switch provides at least some content delivery.
140. The method of claim 133, wherein the processor engine coupling the network switch to a network comprises a network processor.
141. The method of claim 140, further comprising performing look ahead processing within the network processor to off-load processing tasks from the other processor engines.
142. The method of claim 141, wherein the separate processor engines communicate as peers in a peer to peer environment.
143. The method of claim 141, wherein the separate processors engines communicate across a distributed interconnect.
144. The method of claim 143, wherein the distributed interconnect is a switch fabric.

145. The method of 141, wherein the network processor is contained with a network interface engine, the other processing engines comprising a storage processor engine and an application processor engine.

146. The method of claim 145, wherein the network interface engine, the storage processor engine and the application processor engine communicate as peers in a peer to peer environment.

147. The method of claim 146, further comprising performing at least some system management functions in a system management processor engine.

148. The method of claim 147, further comprising tracking system performance within the system management processor engine.

149. The method of claim 147, further comprising implementing system policies with the system management processor engine.

150. The method of claim 133, the network switch providing at least some content delivery.

151. A method of providing a network switching system through the use of a network connectable computing system, comprising:

- providing a plurality of separate processor engines, the processor engines being assigned separate tasks in an asymmetrical multi-processor configuration;
- providing a storage processor engine, the storage processor engine being one of the plurality of separate processor engines;
- providing a network interface connection to at least one of the processor engines to couple the network switching system to a network;
- providing a storage interface connection to the storage processor engine to couple the storage processor engine to a content storage system; and
- accelerating the transmission of data through the network switching system.

152. The method of claim 151, wherein the separate processor engines and the storage processor engine communicate as peers in a peer to peer environment.
153. The method of claim 152, wherein the separate processors and the storage processor engine communicate across a distributed interconnect.
154. The method of claim 153, wherein the distributed interconnect is a switch fabric.
155. The method of claim 153, wherein the processor engine coupling the network endpoint system to a network comprises a network processor.
156. The method of claim 155, further comprising performing look ahead processing within the network processor to off-load processing tasks from the other processor engines.
157. The method of claim 151, wherein the separate processor engine coupling the network endpoint system to a network interface processor engine comprises a network processor.
158. The method of claim 157, further comprising performing look ahead processing within the network processor to off-load processing tasks from the other processor engines.
159. The method of claim 158, wherein the separate processor engines and the storage processor engine communicate as peers in a peer to peer environment.
160. The method of claim 159, wherein the separate processor engines and the storage processor engine communicate across a distributed interconnect.
161. The method of claim 160, wherein the distributed interconnect is a switch fabric.
162. The method of claim 159, wherein one of the separate processor engines is an application processor engine.

163. The method of claim 162, wherein the network interface engine, the storage processor engine and the application processor engine communicate as peers in a peer to peer environment.

164. The method of claim 159, further comprising performing at least some system management functions in a system management processor engine.

165. The method of claim 164, further comprising tracking system performance within the system management processor engine.

166. The method of claim 165, further comprising implementing system policies with the system management processor engine.

167. A network connectable switching system, comprising:
a first processor engine;
a second processor engine, the second processor engine being assigned types of tasks different from the types of tasks assigned to the first processor engine;
a third processor engine, the third processor engine being assigned types of tasks that are different from the types of tasks assigned to the first and second processor engines; and
a distributed interconnection coupled to the first, second and third processor engines, the tasks of the first, second and third processor engines being assigned such that the system operates in staged pipeline manner through the distributed interconnection to perform a switching function.

168. The system of claim 167, wherein the system performs at least some content delivery.

169. The system of claim 167, wherein the first processor engine is a network interface engine comprising a network processor.

170. The system of claim 169, wherein the second processor engine is an application processor engine and the third processor engine is a storage processor engine.

171. The system of claim 170, wherein at least one of the first, second or third processor engines comprises multiple processor modules operating in parallel.

172. The system of claim 171, wherein the application processor engine comprises multiple processor modules operating in parallel and the storage processor engine comprises multiple processor modules operating in parallel.

173. The system of claim 172, wherein the network interface processor engine, the application processor engine, and the storage processor engine communicate in a peer to peer fashion.

174. The system of claim 173, wherein the distributed interconnect is a switch fabric.

175. The system of claim 167, wherein the distributed interconnect is a switch fabric.

176. The system of claim 175, wherein the second processor engine is an application processor engine and the third processor engine is a storage processor engine.

177. The system of claim 176, wherein at least one of the first, second or third processor engines comprises multiple processor modules operating in parallel.

178. The system of claim 177, wherein the application processor engine comprises multiple processor modules operating in parallel and the storage processor engine comprises multiple processor modules operating in parallel.

179. The system of claim 177, wherein the first processor engine is a network interface processor engine comprising a network processor.

180. The system of claim 179, wherein the network interface processor engine, the application processor engine, and the storage processor engine communicate in a peer to peer fashion.

181. A network connectable switching system, comprising:

a first processor engine;

a second processor engine, the second processor engine being assigned types of tasks different from the types of tasks assigned to the first processor engine;

a storage processor engine, the storage processor engine being assigned types of tasks that are different from the types of tasks assigned to the first and second processor engines, the storage processor engine being configured to be coupled to a content storage system; and

a distributed interconnection coupled to the first, second and third processor engines, the tasks of the first, second and third processor engines being assigned such that the system operates in staged pipeline manner through the distributed interconnection.

182. The system of claim 181, wherein the first processor engine is a network interface engine comprising a network processor.

183. The system of claim 182, wherein the second processor engine is an application processor engine.

184. The system of claim 182, wherein at least one of the first, second or storage processor engines comprises multiple processor modules operating in parallel.

185. The system of claim 184, wherein the network interface engine comprises multiple processor modules operating in parallel.

186. The system of claim 182, wherein the network interface processor engine, the second processor engine, and the storage processor engine communicate in a peer to peer fashion.

187. The system of claim 186, wherein the distributed interconnect is a switch fabric.
188. The system of claim 181, wherein the distributed interconnect is a switch fabric.
189. The system of claim 188, wherein the second processor engine is an application processor engine and the third processor engine is a storage processor engine.
190. The system of claim 189, wherein at least one of the first, second or storage processor engines comprises multiple processor modules operating in parallel.
191. The system of claim 190, wherein the application processor engine comprises multiple processor modules operating in parallel and the storage processor engine comprises multiple processor modules operating in parallel.
192. The system of claim 190, wherein the first processor engine is a network interface processor engine comprising a network processor.
193. The system of claim 192, wherein the network interface processor engine, the application processor engine, and the storage processor engine communicate in a peer to peer fashion.
194. A network connectable switching system, comprising:
a first processor engine;
a second processor engine, the second processor engine being assigned types of tasks different from the types of tasks assigned to the first processor engine;
a third processor engine, the third processor engine being assigned types of tasks that are different from the types of tasks assigned to the first and second processor engines; and
a distributed interconnection coupled to the first, second and third processor engines, the tasks of the first, second and third processor engines being assigned such that the system operates in staged pipeline manner through the distributed interconnection,

wherein at least one of the first or second processor engines performs system management functions so as to off-load management functions from the other processor engines.

195. The system of claim 194, wherein the first processor engine is a system management processor engine that performs at least some of the off-loaded management functions.

196. The system of claim 194, wherein the first processor engine is a network interface processor engine that performs at least some of the off-loaded management functions.

197. The system of claim 194, wherein the first processor engine is a network interface processor engine that comprises a network processor.

198. The system of claim 197, wherein the system management functions comprise prioritizing data flow through the system.

199. The system of claim 197, wherein the system management functions comprise quality of service functions.

200. The system of claim 197, wherein the system management functions comprise service level agreement functions.

201. The system of claim 197, wherein the system management functions comprise filtering content requests.

202. The system of claim 201, wherein the first processor engine is a system management processor engine that performs the filtering functions.

203. The system of claim 201, wherein the first processor engine is a network interface processor engine that performs the filtering functions, the network interface processor engine comprising a network processor.

204. The system of claim 194, wherein the third processor engine is a storage processor engine.

205. The system of claim 204, wherein the second processor engine is an application processor engine, wherein at least some system management functions are off-loaded from both the storage processor engine and the application processor engine.

206. The system of claim 194, wherein the system management functions comprise prioritizing data flow through the system.

207. The system of claim 194, wherein the system management functions comprise quality of service functions.

208. The system of claim 194, wherein the system management functions comprise service level agreement functions.

209. The system of claim 194, wherein the system management functions comprise filtering content requests.


210. The system of claim 209, wherein the first processor engine is a system management processor engine that performs the filtering functions.

211. The system of claim 209, wherein the first processor engine is a network interface processor engine that performs the filtering functions, the network interface processor engine comprising a network processor.--

REMARKS

Should any fees under 37 CFR 1.16-1.21 be required for any reason relating to the enclosed materials, the Commissioner is authorized to deduct such fees from Deposit Account No. 10-1205/SURG:157. The examiner is invited to contact the undersigned at the phone number indicated below with any questions or comments, or to otherwise facilitate expeditious and compact prosecution of the application.

Respectfully submitted,



Richard D. Egan
Registration No. 36,788
Attorney for Applicant

O'KEEFE, EGAN & PETERMAN
1101 Capital of Texas Highway South
Building C, Suite 200
Austin, Texas 78746
(512) 347-1611
FAX: (512) 347-1615

APPENDIX
MARKED UP VERSION OF AMENDMENTS
AS REQUIRED BY RULE 121

In The Specification:

On page 1, please replace the paragraph from lines 1-2 with:

--Patent Application for

“NETWORK CONNECTED COMPUTING SYSTEM INCLUDING NETWORK
SWITCH” --

On page 1, please replace the paragraph from lines 6-22 with:

This application is a continuation of Application Serial No. 09/797,413, which was filed March 1, 2001 and is entitled “NETWORK CONNECTED COMPUTING SYSTEM”, which in turn [This application] claims priority from Provisional Application Serial No. 60/187,211, which was filed March 3[0], 2000 and is entitled “SYSTEM AND APPARATUS FOR INCREASING FILE SERVER BANDWIDTH,” from Provisional Application Serial No. 60/246,343, which was filed November 7, 2000 and is entitled “NETWORK CONTENT DELIVERY SYSTEM WITH PEER TO PEER PROCESSING COMPONENTS,” from Provisional Application Serial No. 60/246,335, which was filed November 7, 2000 and is entitled “NETWORK SECURITY ACCELERATOR,” from Provisional Application Serial No. 60/246,443, which was filed November 7, 2000 and is entitled “METHODS AND SYSTEMS FOR THE ORDER SERIALIZATION OF INFORMATION IN A NETWORK PROCESSING ENVIRONMENT,” from Provisional Application Serial No. 60/246,373, which was filed November 7, 2000 and is entitled “INTERPROCESS COMMUNICATIONS WITHIN A NETWORK NODE USING SWITCH FABRIC,” from Provisional Application Serial No. 60/246,444, which was filed November 7, 2000 and is entitled “NETWORK TRANSPORT ACCELERATOR,” and from Provisional Application Serial No. 60/246,372, which was filed November 7, 2000 and is entitled “SINGLE CHASSIS NETWORK ENDPOINT SYSTEM WITH NETWORK PROCESSOR FOR LOAD BALANCING,” the disclosures of each being incorporated herein by reference.

Please replace the paragraph on page 16 from lines 10 -31 and on page 17 from lines 1-3 with the following:

--The processing cores of a network processor are typically accompanied by special purpose cores that perform specific tasks, such as fabric interfacing, table lookup, queue management, and buffer management. Network processors typically have their memory management optimized for data movement, and have multiple I/O and memory buses. The programming capability of network processors permit them to be programmed for a variety of tasks, such as load balancing, network protocol processing, network security policies, and QoS/CoS support. These tasks can be tasks that would otherwise be performed by another processor. For example, TCP/IP processing may be performed by a network processor at the front end of an endpoint system. Another type of processing that could be offloaded is execution of network security policies or protocols. A network processor could also be used for load balancing. Network processors used in this manner can be referred to as "network accelerators" because their front end "look ahead" processing can vastly increase network response speeds. Network processors perform look ahead processing by operating at the front end of the network endpoint to process network packets in order to reduce the workload placed upon the remaining endpoint resources. Various uses of network accelerators are described in the following concurrently filed U.S. patent applications: no. []09/797,412, entitled "Network Transport Accelerator," by Bailey et. al; no. []09/797,507, entitled "Single Chassis Network Endpoint System With Network Processor For Load Balancing," by Richter et. al; and no. []09/797,411, entitled "Network Security Accelerator," by Canion et. al; the disclosures of which are all incorporated herein by reference. When utilizing network processors in an endpoint environment it may be advantageous to utilize techniques for order serialization of information, such as for example, as disclosed in concurrently filed U.S. patent application no. []09/797,197, entitled "Methods and Systems For The Order Serialization Of Information In A Network Processing Environment," by Richter et. al, the disclosure of which is incorporated herein by reference.--

Please replace the paragraph on page 23 from lines 12 - 25 with the following:

--Storage management engine 1040 may employ any suitable method for caching data, including simple computational caching algorithms such as random removal (RR), first-in first-out (FIFO), predictive read-ahead, over buffering, etc. algorithms. Other suitable caching algorithms include those that consider one or more factors in the manipulation of content stored within the cache memory, or which employ multi-level ordering, key based ordering or function based calculation for replacement. In one embodiment, storage management engine may implement a layered multiple LRU (LMLRU) algorithm that uses an integrated block/buffer management structure including at least two layers of a configurable number of multiple LRU queues and a two-dimensional positioning algorithm for data blocks in the memory to reflect the relative priorities of a data block in the memory in terms of both recency and frequency. Such a caching algorithm is described in further detail in concurrently filed U.S. patent application no. []09/797,198, entitled "Systems and Methods for Management of Memory" by Qiu et. al, the disclosure of which is incorporated herein by reference.--

Please replace the paragraph on page 23 from lines 27 - 31 and on page 24 from lines 1 - 5 with the following:

--For increasing delivery efficiency of continuous content, such as streaming multimedia content, storage management engine 1040 may employ caching algorithms that consider the dynamic characteristics of continuous content. Suitable examples include, but are not limited to, interval caching algorithms. In one embodiment, improved caching performance of continuous content may be achieved using an LMLRU caching algorithm that weighs ongoing viewer cache value versus the dynamic time-size cost of maintaining particular content in cache memory. Such a caching algorithm is described in further detail in concurrently filed U.S. patent application no. []09/797,201, entitled "Systems and Methods for Management of Memory in Information Delivery Environments" by Qiu et. al, the disclosure of which is incorporated herein by reference.--

Please replace the paragraph on page 46 from lines 10 - 16 with the following:

--Information gathered by monitoring agents 245 may be employed for a wide variety of purposes including for billing of individual content suppliers and/or users for pro-rata use of one or more resources, resource use analysis and optimization, resource health alarms, *etc.* In addition, monitoring agents may be employed to enable the deterministic delivery of content by system 200 as described in concurrently filed, co-pending United States patent application number []09/797,200, entitled “Systems and Methods for the Deterministic [Delivery of Data and Services]Management of Information,” which is incorporated herein by reference.--

Please replace the paragraph on page 28 from lines 16 - 26 with the following:

--In yet other embodiments, the processing modules may be specialized to specific applications, for example, for processing and delivering HTTP content, processing and delivering RTSP content, or other applications. For example, in such an embodiment an application processing module 1070a and storage processing module 10[5]40a may be specially programmed for processing a first type of request received from a network. In the same system, application processing module 1070b and storage processing module 10[5]40b may be specially programmed to handle a second type of request different from the first type. Routing of requests to the appropriate respective application and/or storage modules may be accomplished using a distributive interconnect and may be controlled by transport and/or interface processing modules as requests are received and processed by these modules using policies set by the system management engine.--

Please replace the paragraph on page 34 from lines 2 - 13 with the following:

--In yet another embodiment, at least two network interface modules 1030a and 1030b may be provided, as illustrated in FIG. 1A. In this embodiment, a first network interface engine 1030a may receive incoming data from a network and pass the data directly to the second

network interface engine 1030b for transport back out to the same or different network. For example, in the remote or live broadcast application described above, first network interface engine 1030a may receive content, and second network interface engine 1030b provide the content to the network 1020 to fulfill requests from one or more clients for this content. Peer-to-peer level communication between the two network interface engines allows first network interface engine 1030a to send the content directly to second network interface engine 1030b via distributive interconnect 1080. If necessary, the content may also be routed through transport processing engine 1050, or through [network]transport processing engine 1050 and application processing engine 1070, in a manner described above.--

Please replace the paragraph on page 34 from lines 23 - 31 with the following:

--The content delivery system 1010 described above is configured in a peer-to-peer manner that allows the various engines and modules to communicate with each other directly as peers through the distributed interconnect. This is contrasted with a traditional server architecture in which there is a main CPU. Furthermore unlike the arbitrated bus of traditional servers, the distributed interconnect 1080 provides a switching means which is not arbitrated and allows multiple simultaneous communications between the various peers. The data and communication flow may by-pass unnecessary peers such as the return of data from the storage management processing engine 10[6]40 directly to the network interface processing engine 1030 as described with reference to FIG. 1B.--

Please replace the paragraph on page 42 from lines 23 - 31 with the following:

--It will be understood with benefit of this disclosure that the hardware embodiment and multiple engine configurations thereof illustrated in FIGS. 1C-1F are exemplary only, and that other hardware embodiments and engine configurations thereof are also possible. It will further be understood that in addition to changing the assignments of individual processing modules to particular processing engines, distributive interconnect 1080 enables the [vary]various processing flow paths between individual modules employed in a particular engine configuration in a

manner as described in relation to FIG. 1B. Thus, for any given hardware embodiment and processing engine configuration, a number of different processing flow paths may be employed so as to optimize system performance to suit the needs of particular system applications.--

Please replace the paragraph on page 44 from lines 15 - 28 with the following:

--FIG. 2 illustrates a single networking subsystem processor module 205 that is provided to perform the combined functions of network interface processing engine 1030 and transport processing engine 1040 of FIG. 1A. Communication and content delivery between network 260 and networking subsystem processor module 205 are made through network connection 270. For certain applications, the functions of network interface processing engine 1030 and transport processing engine 1050 of FIG. 1A may be so combined into a single module 205 of FIG. 2 in order to reduce the level of communication and data traffic handled by communications path 230 and data movement path 235 (or single switch fabric), without adversely impacting the resources of application processing engine or subsystem module. If such a modification were made to the system of FIG. 1A, content requests may be passed directly from the combined interface/transport engine to network application processing engine 1070 via distributive interconnect 1080. Thus, as previously described the functions of two or more separate content delivery system engines may be combined as desired (e.g., in a single module or in multiple modules of a single processing blade), for example, to achieve advantages in efficiency or cost.--